arXiv:1403.8020v1 [cs.DC] 31 Mar 2014

An Efficient Thread Mapping Strategy for Multiprogramming on Manycore Processors

Ashkan TOUSIMOJARAD, Wim VANDERBAUWHEDE
School of Computing Science
University of Glasgow

Abstract. The emergence of multicore and manycore processors is set to change the parallel computing world. Applications are shifting towards increased parallelism in order to utilise these architectures efficiently. This leads to a situation where every application creates its desirable number of threads, based on its parallel nature and the system resources allowance. Task scheduling in such a multithreaded multiprogramming environment is a significant challenge. In task scheduling, not only the order of the execution, but also the mapping of threads to the execution resources is of a great importance. In this paper we state and discuss some fundamental rules based on results obtained from selected applications of the BOTS benchmarks on the 64-core TILEPro64 processor. We demonstrate how previously efficient mapping policies such as those of the SMP Linux scheduler become inefficient when the number of threads and cores grows. We propose a novel, low-overhead technique, a heuristic based on the amount of time spent by each CPU doing some useful work, to fairly distribute the workloads amongst the cores in a multiprogramming environment. Our novel approach could be implemented as a pragma similar to those in the new task-based OpenMP versions, or can be incorporated as a distributed thread mapping mechanism in future manycore programming frameworks. We show that our thread mapping scheme can outperform the native GNU/Linux thread scheduler in both single-programming and multiprogramming environments.

Keywords. Thread Mapping, Manycore Processors, Multiprogramming, OpenMP, Task Parallelism

Introduction

Recently, task parallelism has gained a great importance in parallel computing. Some task parallel frameworks provide explicit task parallelism [1], while in others parallelism is expressed implicitly [2]. The OpenMP *task* directives can be used to define units of independent work as *task*s. However, the scheduling decisions are leaved to the runtime systems. [3] suggests the improvement of OpenMP runtime support. Nanos v4 [4] is an OpenMP runtime library that provides some mechanisms to allow the user to choose between different task scheduling policies. In [5], two families of the schedulers are added to the Nanos runtime: Breadth-First schedulers and Work-First schedulers. The main focus is on how threads execute the tasks, while our focus in this paper is on where to place those threads.

Task scheduling is the arrangement of tasks in time and space on the available execution resources. Therefore, not only the order of execution, but also optimum mapping of all threads to the existing resources is crucial to get the best result. In [6], binding threads to the physical processors, also referred to as *thread affinity*, is proposed to be a part of the OpenMP standard. It has been shown how a simple Round-Robin mapping scheme can improve the performance.

An optimal solution to task scheduling problem cannot be found in polynomial time which means it is an NP-hard problem. This motivates researchers to develop heuristics to find near optimal solutions [7]. Previous studies show that efficient mapping of threads to specific cores, providing load balancing among the cores would result in improved performance. In [8], a sophisticated scheduling scheme in multiprogramming environment is developed, based on dynamically prediction of the applications' scalability. It motivated us by highlighting this fundamental problem that the performance of some applications without linear scalability tend to decrease when more number of cores are allocated to them. Limiting the number of threads is a way to overcome this issue. Therefore, for an application with smaller number of threads than the number of cores, it becomes important to determine where to map those threads. We will address this issue by calculating the current CPU load of the tiles before mapping the threads to them.

Generally, scheduling policies can be evaluated based on system-oriented or user-oriented criteria [9]. A system-oriented metric is based on the system's perspective and quantifies how effectively and efficiently the system utilises the resources, while the focus of a user-oriented metric is on the behaviour of the system from user's perceptive, e.g. how fast a single program is executed. An example of system-oriented metrics is *throughput*, which is the number of programs completed per unit of time. *Turnaround time* is an example of user-oriented metrics, which is the time between submitting a job and its completion. In this work, we have used *Turnaround time* to evaluate different mapping techniques on the TILEPro64 machine.

1. Selected Benchmarks

We show how different thread mapping strategies can affect the performance of four benchmarks from the Barcelona OpenMP Tasks Suite (BOTS) [10], selected for their different characteristics. The mapping techniques are low-overhead, and can be combined with different cut-off strategies and applied on either *tied* or *untied* tasks.

It is important to note that the applications which do not scale very well are more challenging for parallel computing. Embarrassingly parallel algorithms are easy to parallelise since the tasks are completely (or almost) independent. They can easily run on different processing cores without the need to share data or exchange any information with each other. We have used a benchmark that scales approximately linearly (NQueens), one that does not scale well when the number of threads grows (Strassen), and two others that reach their saturation phases (Sort and Health). The input sets are chosen in such a way that the turnaround times of the programs range from a few seconds to a few tens of seconds. The aim is to show that the overhead of the proposed mapping technique is negligible, even for programs with small turnaround times.

The target platform is the TILEPro64, which runs Tile Linux that is based on the standard open-source Linux version 2.6.26. The C compiler used is the one provided in

the Multicore Development Environment (MDE) 3.0 from Tilera Corporation, which is called Tile-cc and is based on the GCC 4.4.3. The only change made to the BOTS 1.1.2 configuration file is the name of the compiler.

1. **Sort (untied):** Sorts a random permutation of n 32-bit numbers with a fast parallel sorting variation of the ordinary merge sort. First, it divides an array of elements in two halves, sorting each half recursively, and then merging the sorted halves with a parallel divide-and-conquer method rather than the conventional serial merge. Tasks are used for each split and merge. When the array is too small, a serial quick sort is used to increase the task granularity. We have used the default cut-off values (2048) when sorting an array of 50M integers. To avoid the overhead of quick sort, an insertion sort is used for small arrays of 20 elements.
2. **Health (manual-tied):** This program simulates de Columbian Health Care System. Each element in its multilevel lists represents a village with a list of potential patients and one hospital. The status of a patient in the hospital could be waiting, in assessment, in treatment, or waiting for reallocation. Each village is assigned to one task. The probabilities of getting sick, needing a convalescence treatment, or being reallocated to an upper level hospital are considered for the patients. At each time-step, all patients are simulated according to these probabilities. To avoid indeterminism in different levels of the simulation, one seed is used for each village. Therefore, all probabilities computed by a single task are identical across different executions and are independent of all other tasks. three different input sizes are available in the benchmark suite. We have used them in different scenarios. However, the performance scalability of the single program is presented using the medium-size input.
3. **Strassen (tied):** The Strassen algorithm employs a hierarchical decomposition of a matrix for multiplication of large dense matrices. Decomposition is performed by dividing each dimension of the matrix into two parts of equal size. For each decomposition a task is created. A matrix size of 2048×2048 is used for the purposes of this experiment.
4. **NQueens (manual-untied):** The NQueens benchmark computes all solutions of the n-queens problem, whose aim is to find a placement for n queens on an $n \times n$ chessboard such that none of the queens attack any other. It uses a backtracking search algorithm with pruning. A task is created for each step of the solution,and it has an almost linear speed-up.

2. Mapping Strategies

We have performed the experiments with four different mapping strategies. In the first part of the experiments, we measure the execution time of each benchmark under these mapping policies in a single-programming environment. The second part is to investigate the behaviour of these mapping strategies in a multiprogramming environment. In each of our mapping schemes, every thread decides about its mapping itself. It first finds a suitable core, maps itself to it and starts doing some work or goes to sleep. In the OpenMP code, this happens after the *parallel* keyword, which is the point where the thread creation happens.

2.1. Linux Scheduler

The first option is to leave any scheduling decision to the native Linux scheduler. Tilera's version of SMP Linux, called Tile Linux is based on the standard open-source Linux version 2.6.26. The default scheduling strategy in Linux is a priority-based dynamic scheduling that allows for thread migration to idle cores in order to balance the runqueues.

Having a single runqueue for all processors in a Symmetric Multiprocessing (SMP) system, and using a single runqueue lock were some of the drawbacks of the Linux 2.4 scheduler. Linux 2.6 implemented a priority-based scheduler known as the O(1) scheduler, which means the time needed to select the appropriate process and map it to a processor is constant. One runqueue data structure per each processor keeps track of all runnable tasks assigned to that processor.

At each processor, the scheduler picks a task from the highest priority queue. If there are multiple tasks in that queue, they are scheduled in a Round-Robin manner. There is also a mechanism to move tasks from the queues of one processor to those of another. This is done periodically by checking whether the *cpu_load* is imbalanced. In the Linux terminology, *cpu_load* is the average of the current load and the old load. The current load is the number of active tasks in the CPU's runqueue multiplied by `SCHED_LOAD_SCALE`, which is used to increase the resolution of the load [11]. What we will refer to as *load* in our proposed technique is the amount of time spent in each processor doing some useful work.

2.2. Static Mapping

In the static mapping strategy, threads are pinned to the processing cores based on their *thread_id*s in an ordered fashion. The decision is taken at compile time, which would cause an obvious disadvantage: It cannot tune itself with multiprogramming, since every program follows the same rule, and if the number of threads are less than the number of cores, then some cores get no threads at all. It might be discussed why at the first place, the number of threads in each program should be less than the number of cores. The answer to this question can be found in the applications which do not have linear speed-up, and after a certain number of threads reach their saturation phase. An example is the Sort program in this benchmark suite. This behaviour is reported in [8] as well, where the authors emphasise that the scalability of some programs tend to saturate at some points, and their performance is degraded by adding more cores.

2.3. Basic Lowest Load (BLL)

The Lowest Load mapping technique is presented as two different methods. The first one, assumes the term *load* as an equivalent to a thread. Therefore, if one thread is mapped to a core, the core's load becomes 1. We call this method *Basic Lowest Load* (BLL). It fills out the cores of the system in a Round-Robin fashion. Again, this technique is not aware of what is going on in the system. There are many situations in which some idle cores are ignored. We will show an example in the next section.

2.4. The Extended Lowest Load (XLL)

The XLL gets the cores' information from the */proc/stat* file in Linux. The amount of time each core has done different types of work is specified with a number of time units. The time units are expressed in USER_HZ or Jiffies, which are typically hundredths of a second. The number of Jiffies in user mode is selected as *load*. In this technique, every thread scans the current *load*s of the cores. It then searches for a core with the least change from its old *load* value. The thread maps itself to that core and starts working. In other words, the actual target of this policy is the least busy core. Except from its dynamic awareness of the system, another difference with the BLL becomes highlighted when a thread is created but goes instantly to the sleeping mode. The XLL automatically finds the sleeping threads since they do not produce any *load*, and hence more threads can be assigned to the corresponding cores, while the BLL only counts the number of pinned threads to the core, no matter if they are sleeping or doing some work. The algorithm for the XLL methodology is as follows:

Algorithm 1 The XLL Methodology

```
1:  procedure FINDBESTTARGET
2:      GetTheLock();
3:      for each int i in Cores  do
4:          Scan(CurrentLoad[i]);                          ▷ Scans from the /proc/stat file
5:          Cores[i].change = CurrentLoad[i] - Cores[i].load + Cores[i].pinned;
6:          Cores[i].load = CurrentLoad[i] + 10;           ▷ Creates a better resolution
7:      end for
8:      for each int i in Cores  do
9:          if Cores[i].change <Cores[BestTarget].change then
10:             BestTarget = i;                            ▷ Finds the least busy core
11:         end if
12:     end for
13:     SetAffinity(BestTarget);
14:     Cores[BestTarget].pinned++;          ▷ Increments the number of pinned threads
15:     ReleaseTheLock();
16: end procedure
```

The proposed methodology requires a globally shared data structure that keeps track of the system's cores. This data structure can be implemented in a runtime system as in our work, or can be embedded in the Linux kernel. It is worth mentioning that this methodology is portable across similar multicore/manycore platforms.

3. Results

The first step is to analyse the performance behaviour of each benchmark individually. Figure 1 shows how our 4 applications from the BOTS benchmark suite scale on the TILEPro64. As mentioned before, not all the applications benefit from the increased number of threads. Thus, in order to get the best performance, one solution is to limit the number of threads based on the scalability of the application.

For a single-program workload, The Static, BLL, and XLL techniques behave very similarly. The comparison with the Static mapping verifies that the overhead of the BLL and XLL techniques are negligible. Moreover, we can see that even for a single-program workload, the proposed mapping technique (XLL) works better than Linux.

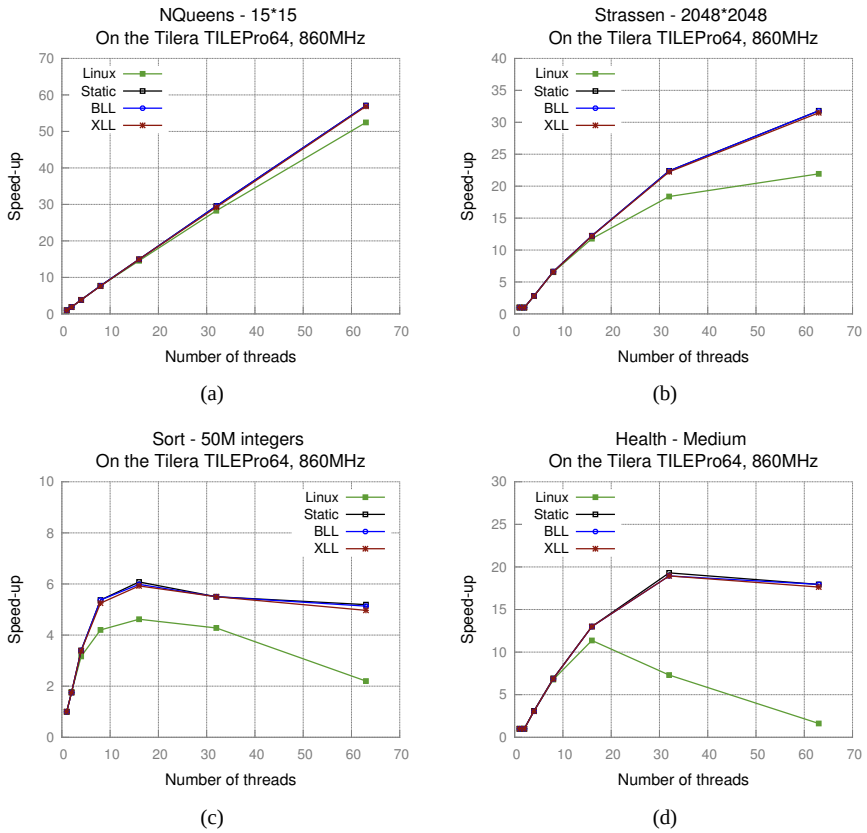

Figure 1. Speed-up of the selected benchmarks under different mapping policies (a) NQueens 15×15 board, (b) Strassen 2048×2048 matrix, (c) Sort 50M integers, (d) Health Medium input

For multiprogram workloads, we have considered three different scenarios to show how the XLL mapper can result in better performance. It is worth mentioning that each experiment was run 20 times.

First, we have to show why BLL, which is a simple Round-Robin mapping algorithm is inefficient. For this purpose, we have considered three Health programs, each of which with 32 threads (that gave us the best performance in the single-program scenario). Two programs have large inputs and one has a small input. The programs enter the system with the interval of 6 Secs. We have previously discussed why Static mapper cannot handle multiprogramming scenarios. The inefficiency of the BLL is also evident from Figure 2.

The first scenario clearly shows that the XLL is the winning policy. The scenario is designed in such a way that the program with the small input data set finishes before the second large program enters the system. In the case of BLL, the threads of the first

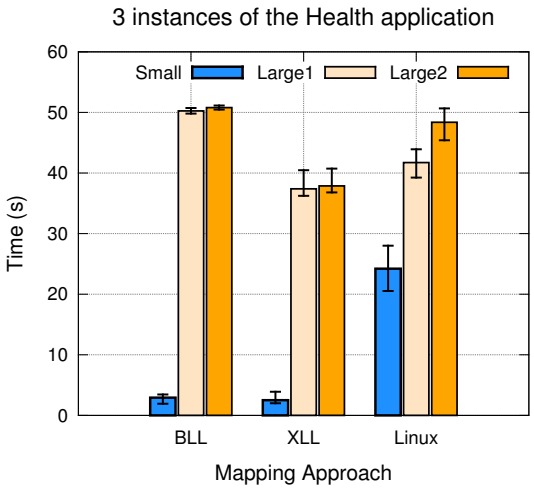

Figure 2. The first scenario: The inefficiency of the BLL

large program are mapped to the first 32 cores of the system. The threads of the small program are mapped to the last 31 cores plus the first core (there are 63 cores to use). Then the small program finishes and the second large program enters the system, but the BLL cannot use the recently freed cores. Instead, based on the Round-Robin algorithm, the threads of the second large program are mapped to the cores 2 to 33, while most of these cores (except one of them) are already busy serving the first large program.

The second scenario is to run all four programs selected from the BOTS at the same time. Although they start at the same time, their thread creation time is different. This is due to the fact their initialisation phases and memory allocation times are different. Recall that thread creation happens whenever the execution reaches the *parallel* keyword in the OpenMP code. According to the single-program performance, the Sort program (50M integers) is limited to 16 threads, the Health program (Medium input) is limited to 32 threads, and both Strassen (2048×2048) and NQueens (15×15) are run with 63 threads. The result is shown in Figure 3.

Once again, the XLL results in better performance. The turnaround times for all 4 programs are smaller when the XLL policy is applied.

The third scenario gives a better insight on how the XLL mapper outperforms the Linux scheduler when the system is busy. For this scenario, we have used 10 identical instances of the Sort program arriving the system one after the other with the interval of 1 second. The result is depicted in Figure 4.

Figure 4 shows that the results with both policies are significantly better when each Sort program uses 16 threads rather than 63. It again verifies that increasing the number of threads does not necessarily result in better performance. It is also evident how much our novel XLL mapping technique can outperform the native Linux scheduler in a multiprogramming environment.

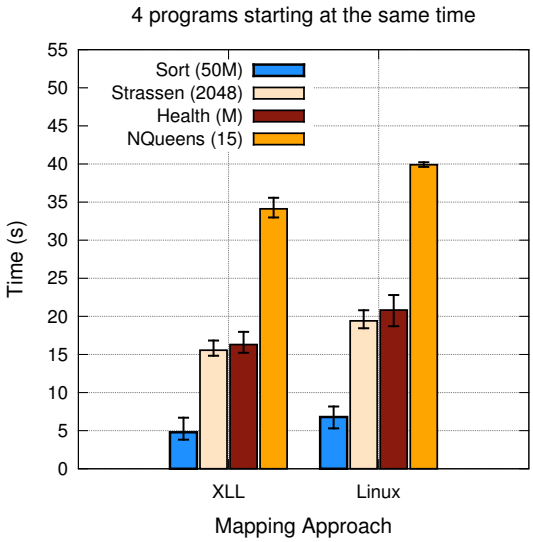

Figure 3. The second scenario: Running selected programs as the same time

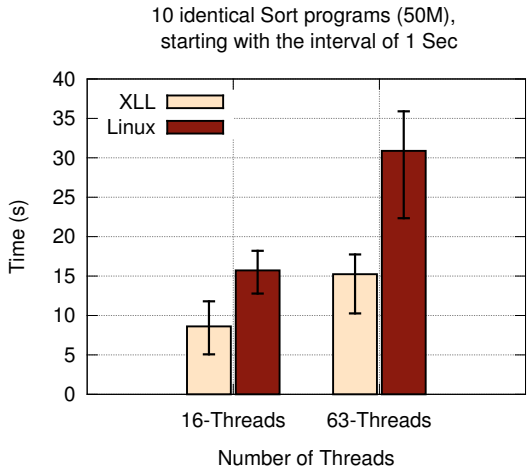

Figure 4. The third scenario: Running 10 identical instances of the Sort program

4. Conclusion

In this work, we have performed an analysis of multiprogramming performance on a Tilera manycore system using the BOTS benchmark. We have shown that although the current SMP Linux scheduler performs well for small numbers of threads (up to 8 threads), the scheduling performance degrades for increasing numbers of cores and threads, leading to poor performance on manycore systems. We observe that increasing the number of threads does not necessarily lead to better performance and can even de-

grade the overall performance. With a smaller number of threads than the number cores, different mapping configurations are possible, which allows to optimise the performance.

We have presented a novel, low-overhead, fundamental mapping strategy to provide load balancing in a multithreaded multiprogramming environment: the *Extended Lowest Load* (XLL) technique uses a heuristic to find the optimal target core for each thread. Although this work is in an early stage, the results are very promising. In this paper, we have shown how our XLL mapping scheme outperforms the native GNU/Linux thread scheduler in different scenarios.

References

[1] Eduard Ayguadé, Nawal Copty, Alejandro Duran, Jay Hoeflinger, Yuan Lin, Federico Massaioli, Xavier Teruel, Priya Unnikrishnan, and Guansong Zhang. The design of openmp tasks. *Parallel and Distributed Systems, IEEE Transactions on*, 20(3):404–418, 2009.

[2] Ashkan Tousimojarad. A parallel task composition approach to manycore programming. *PLACES 2013*, page 29, 2013.

[3] Stephen L Olivier and Jan F Prins. Comparison of openmp 3.0 and other task parallel frameworks on unbalanced task graphs. *International Journal of Parallel Programming*, 38(5):341–360, 2010.

[4] Xavier Teruel, Xavier Martorell, Alejandro Duran, Roger Ferrer, and Eduard Ayguadé. Support for openmp tasks in nanos v4. In *Proceedings of the 2007 conference of the center for advanced studies on Collaborative research*, pages 256–259. IBM Corp., 2007.

[5] Alejandro Duran, Julita Corbalán, and Eduard Ayguadé. Evaluation of openmp task scheduling strategies. In *OpenMP in a New Era of Parallelism*, pages 100–110. Springer, 2008.

[6] Guansong Zhang. Extending the openmp standard for thread mapping and grouping. In *OpenMP Shared Memory Parallel Programming*, pages 435–446. Springer, 2008.

[7] Oliver Sinnen. *Task scheduling for parallel systems*, volume 60. Wiley-Interscience, 2007.

[8] Hiroshi Sasaki, Teruo Tanimoto, Koji Inoue, and Hiroshi Nakamura. Scalability-based manycore partitioning. In *Proceedings of the 21st international conference on Parallel architectures and compilation techniques*, pages 107–116. ACM, 2012.

[9] William Stallings. *Operating Systems, Internals and Design Principles, 7th Edition*. Pearson: Prentice Hall, 2012.

[10] Alejandro Duran, Xavier Teruel, Roger Ferrer, Xavier Martorell, and Eduard Ayguade. Barcelona openmp tasks suite: A set of benchmarks targeting the exploitation of task parallelism in openmp. In *Parallel Processing, 2009. ICPP'09. International Conference on*, pages 124–131. IEEE, 2009.

[11] Josh Aas. Understanding the linux 2.6.8.1 cpu scheduler. *Retrieved Oct*, 16:1–38, 2005.